



INCITE

InterNet Control and Inference Tools at the Edge


R. Baraniuk R. Nowak E. Knightly R. Riedi
M. Coates X. Wang V. Ribeiro S. Sarvotham

NMS PI meeting Atlanta October 2001

Effort 1

Chirp Probing

Objective: Reduced complexity, multiscale
link models with known accuracy




Innovative Ideas

Multifractal analysis
Multiplicative modeling
Multiscale queuing
Chirps for probing

Impact

Congestion control
Workload balancing at servers
Dynamical streaming
Pricing on connection basis

Chirp Probe Cross-Traffic Inference



New Ideas

Probing multiple hops

Network calculus

Probe size distributions

Probing buffer at core router

Passive inference (Sprint)

Tech Transfer

CAIDA (chirping as a monitoring tool)

Stanford (SLAC) (chirps and PingER)

Los Alamos (LANL)

Sprint Labs

Microsoft (streaming applications)


UCRiverside (expertise on self-similarity)

Effort 2

Connection-level Analysis and Modeling of Network Traffic

Aggregate Statistics



- Aggregate load on link
 - Time stamped headers
- Positive process
- Burstiness
 - LRD (large scale)
 - Non-Gaussian (small scale)



Objective : Origins of small scale bursts


Bursts in the ON/OFF framework

- ON/OFF model
 - Superposition of sources
 - Connection level model
- Explains large scale variability:
 - LRD, Gaussian
 - Cause: Costumers
 - Heavy tailed file sizes !!
- Small scale bursts:
 - Non-Gaussianity
 - Conspiracy of sources ??
 - Flash crowds ??
(dramatic increase of active sources)




Non-Gaussianity: A Conspiracy?

Load: Bytes per 500 ms



- The number of active connections is close to Gaussian; provides no indication of bursts in the load

Number of active connections





- Indication for:
 - **No conspiracy** of sources
 - No flash crowds

Non-Gaussianity: a case study

Typical **Gaussian** arrival
(500 ms time slot)

Histogram of load
offered in same time
bin per connection:


Considerable **balanced**
"field" of connections



Typical **bursty** arrival
(500 ms time slot)

Histogram of load
offered in same time
bin per connection:


**One connection
dominates**



Non-Gaussianity and Dominance

Circled in Red:

Instances where one connection contributes over 50% of load (resolution 500 ms)




- Dominant connections correlate with bursts


Non-Gaussianity and Dominance

Systematic study: **time series** separation

- For each bin of 500 ms:
remove packets of the **ONE strongest** connection
- Leaves “Gaussian” residual traffic




Overall traffic



1 Strongest connection

+



Residual traffic

Separation on Connection Level

Definition:

- **Alpha** connections:

Peak rate $>$ mean arrival rate + 1 std dev

- **Beta** connections: Residual traffic

- Findings are similar for

- Auckland (2000+2001), Berkeley, Bellcore, DEC
- 500ms, 50ms, 5ms resolution

Alpha Traffic Component

- There are **few Alpha** connections
 - < 1% (AUCK 2000: 427 of 64,087 connections)
 - 3% of load

- Alpha connections cause **bursts**:

Alpha is extremely bursty


Beta is little bursty

Overall traffic is quite bursty

Balanced (50% alpha) very bursty


Multifractal spectrum:

Wide spectrum means bursty




Beta Traffic Component

- Constitutes main load
- Governs **LRD** properties of overall traffic
- Is **Gaussian** at sufficient utilization (Kurtosis = 3)
- Is well matched by ON/OFF model



Beta traffic



Number of connections = ON/OFF

What Causes Alpha Connections?

- Potential causes:
 - TCP slow-start peculiarities
 - Start/End of “massive” flows
 - Re-routing
 - Heterogeneity in bandwidths
- Look for: systematic explanation
 - First two: anywhere in network
 - Last two: locality in network important



Origins of Alpha Traffic 1

- Observation 1:
Alpha connections **cluster** into e2e groups
 - **e2e group**:
connections with same source-receiver pair
 - 85 (out of total 6960) e2e groups contain at least one of the 427 alpha connections (AUCK)
 - **Locality matters**
- Excludes **TCP slow start** and **start/end** of “massive” flows as systematic causes

Origin of Alpha Traffic 2




- Observation 2: If one connection in e2e-group is alpha, then all connections are
 - **Unlimited** (Peak rate $> \frac{1}{2}$ Total transfer)
 - TCP control mechanism does not become effective
 - and/or
 - **Alpha** (Peak $>$ Threshold)
 - Causes burst

Peak rate and total load per connection for two e2e groups



Origin of Beta traffic

- Observation 3: If **no** connection in e2e-group is **alpha**, then all connections are
 - **Limited** by same bottleneck bandwidth



Evidence for: **bandwidth** matters

Confirms **ON/OFF** as a good model of Beta traffic

Simple Connection Taxonomy

		connection speed	
		slow	fast
file size	small	Beta	Beta
	large	Beta	Alpha

Bursts arise from large transfers over fast links.

*This is the **only** systematic reason*

Modeling Network Traffic

Physical Model

- Traffic (user): superposition of ON/OFF sources requesting files with heavy tailed size
- Network: heterogeneous bandwidth
 - variable sending-rates (fixed per ON/OFF source)
- Explains properties of traffic:
 - LRD: heavy tailed transfer of beta sources (crowd)
 - Bursts: few large transfers of few alpha sources


Mathematical Model

- Traffic = Alpha + fractional Gaussian noise


Impact 1: Simulation

- ns: **topology** should include a few **alpha** links

Simple: equal bandwidth




Realistic: **heterogeneous** end-to-end bandwidth



Impact 2: Queueing

- **Beta:** rules small queues
- **Alpha:** rules the rare extremely long queues
- Needs theoretical work



Summary

- Connection level analysis of “all” available traces
- Typically **one dominant** connection during burst
- **Alpha** traffic (peak rate > burst threshold)
 - Few connections. Responsible for **bursts**
 - Origin: Large transfer over **high bandwidth** paths
 - Bursts are less pronounced at high utilization
- **Beta** traffic (residual):
 - Main load. Responsible for **LRD**
 - Origin: Crowd with **limited** bandwidth
 - **Gaussian** at sufficient utilization

Future work

- Queueing analysis
- Mathematical model for Alpha
- Further verification using
 - More (new) traces
 - Simulation (ns, testbed) → Realism
 - Monitoring of real network
- Influence of other parameters on presence of bursts
 - Utilization
 - Delay
 - Topology

incite.rice.edu